

Hands-On Oracle® Database 10g: Backup, Recovery and Server Tuning - 5 Days

Course 518 Overview

- You Will Learn How To**
- Deploy backup and recovery strategies to safeguard Oracle 10g databases
 - Perform recovery operations to maintain consistent and available data
 - Failover to hot standby databases for high availability
 - Exploit automated backup and recovery techniques with Recovery Manager
 - Implement flashback and logmining to roll back user and logical errors
 - Tune the Oracle 10g server by controlling memory allocation, disk I/O and locking/latch contention
- Course Benefits** The guaranteed continuous operation of databases is a critical requirement of any successful organization. Oracle Database 10g includes an increased range of features to ensure optimal performance, plus protection and recovery from failure. In this course, you apply techniques to solve a range of issues, including media and site failures, invalid user input, lock contention, and poor memory and storage allocation.
- Who Should Attend** Those interested in safeguarding Oracle 10g databases. Knowledge of database administration at the level of Course 594, "Oracle Database 10g Administration," is assumed.
- Hands-On Training** You gain extensive experience using Oracle 10g. Exercises include:
- Implementing a backup strategy with Recovery Manager
 - Recovering the database from serious errors
 - Enabling fast recovery by rolling backup files forward
 - Transporting tablespaces across databases
 - Setting up LogMiner to recover bad transactions
 - Creating hot standby databases
 - Optimizing query performance
 - Controlling shared pool usage
 - Diagnosing lock contention with Database Control 10g

Hands-On Oracle® Database 10g: Backup, Recovery and Server Tuning - 5 Days

Course 518 Outline

The Oracle 10g Architecture

- Instances and databases
- The System Global Area (SGA) and background processes
- Automated storage configurations

Implementing Backup and Recovery Strategies

Protecting the database

- Choosing appropriate backup strategies
- Identifying points of possible failure
- Safeguarding redo logs and control files
- Managing archive logs

Performing backups

- Generating full database backups
- Implementing partial online backups
- Backing up and creating control files

Transporting tablespaces

- Establishing a consistent dataset
- Migrating across hardware platforms

Achieving full database recovery

- Recovering tablespaces and datafiles
- Reconstructing datafiles without backups
- Repairing online redo logs
- Flashing back the entire database

Disaster Recovery Using Hot Standby Databases

Configuring the standby environment

- Creating and synchronizing the standby database
- Maintaining a read-only standby database
- Preserving the performance of the primary database
- Physical vs. logical standby databases

Moving operations to the standby

- Failing over and switching to the standby facility
- Achieving No-Data-Loss recovery

Automating Backup and Recovery with Recovery Manager (RMAN)

Setting up Oracle Recovery Manager

- Creating the Recovery Catalog
- Registering databases for recovery
- Configuring channels and redundancy

Recovery Manager backups

- Full and incremental backups

- Scripting the backup activity
- Maintaining the catalogue
- Listing and reporting on backups

Performing automated recovery

- Restoring and recovering data from backup sets and image copies
- Rolling forward image copies with incremental backups
- Tuning backup processes with block change tracking

Analyzing Redo Log Files with LogMiner

Configuring LogMiner

- Creating the LogMiner dictionary
- Manually listing log files to be analyzed
- Automatically choosing log files using continuous mining
- Adding supplemental log data

Interpreting LogMiner information

- Auditing changes to specific columns
- Performing fine-grained recovery from logical corruptions

Tuning the Oracle 10g Server

Tuning the SGA

- Maximizing the use of the shared pool
- Tuning the buffer cache with the Buffer Cache Advisor
- Reducing I/O with multiple buffer pools
- Monitoring latch contention
- Configuring Startup with triggers

Automating memory management

- Managing Program Global Area (PGA)
- Enabling dynamic memory allocation

Optimizing transactions and queries

- Determining block contention
- Compressing table data
- Monitoring latch contention and waits

Applying Diagnostic Techniques

- Managing the workload repository
- Identifying contention with the Automatic Database Diagnostic Monitor (ADDM)
- Examining the alert log